


Exceptional service in the national interest



Data Transfer Tools - Exascale Performance Requirements

Susie McRee (Project Lead), Bill Collins (Production Lead)

Joint Facilities User Forum on Data-Intensive Computing



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND Number: 2014-4357C.

Agenda

- Introduction
- Current
 - Technology
 - Tools
 - Performance
- Towards Exascale
 - Requirements
 - Futures
 - Approach

- High Performance Storage System collaboration
 - Executive/Technical member
 - Transfer tool development
 - Parallel FTP – HPSS, intra-host (PFTP server)
 - Transfer Agent (TA) – SNL & GEL (Gleicher Enterprises Ltd.)
 - HSI/HTAR – development, support
- Tri-Lab (SNL, LLNL, LANL)
 - Tool installation, configuring, tuning
 - WAN (DisCom)
- Advanced Computing at Extreme Scale(ACES)
 - Cielo

Current, Tri-lab Technologies

- Network Technology
 - 10Gb LAN, near future 40Gb (?)
 - InfiniBand
 - 10Gb Tri-Lab WAN, Future 100Gb (SNL, LANL)
- File system
 - Lustre
 - Panasas
 - Traditional
 - Network File System (NFS), Extended File system (EXT), etc.

Tri-Lab tools

- Parallel File Transport (PFTP)
 - Primary HPSS development, SNL
 - HPSS, client, (host) server
 - Multi-path, Multi-network
- Transfer Agent (TA)
 - Developed by SNL, GEL
 - Multi-host – large file performance
- HSI/HTAR
 - Multi-thread, concurrent transfers
 - HSI, Integrated TA – contracted, GEL
 - Host server (aka non-HPSS HSI gateway) – contracted, GEL

Tri-lab performance

- “Large” file transfers
 - 2 X 2TB, LANL (FTA) to SNL (HPSS) - 630MB/s
 - HSI, two simultaneous transfers
 - Pdata protocol (request/send)
 - 2.6ms RTT
 - 2 X 2TB LLNL (HPSS) to LANL (FTA) – 195MB/s ea.
 - 40.1ms RTT
- Many/small file transfers
 - 4 X HTAR each 135K files, 4.5TB
 - LANL (FTA) to SNL HPSS, aggregate 534MB/s
- Largest SNL HPSS file, 20TB
 - Who’s got the 50TB?

Exascale demands

- Projected data demands
 - 2-6PB, 10's to a few 100's million files
 - 10 -> 100Gb/s WANS
 - Same processing/storage use cases(?)
- Broad requirements
 - Discussions with LLNL, LANL, ORNL, NERSC
 - Still gathering requirements – wpcolli@sandia.gov,
srmcree@sandia.gov

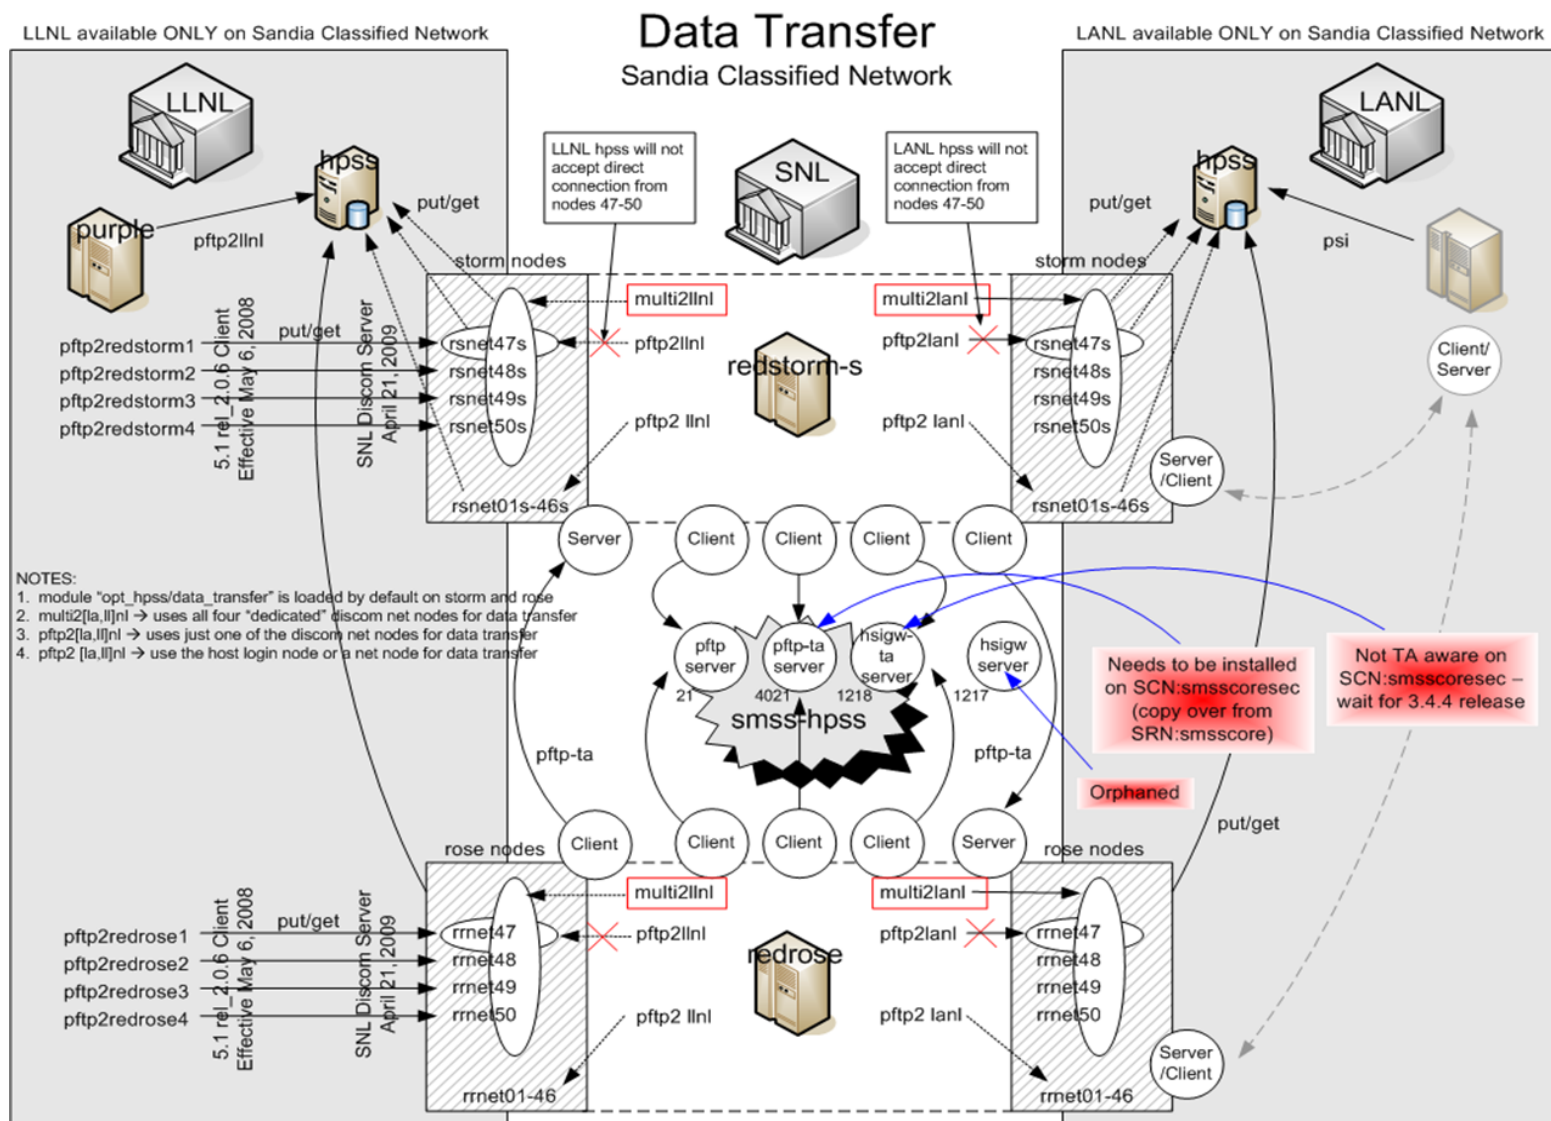
Exascale desirables

- User
 - Initiate transfers
 - Monitor transfers
- Presentation
 - GUI – ease of use
 - Web based
 - Cell phone app (?)

Customer view

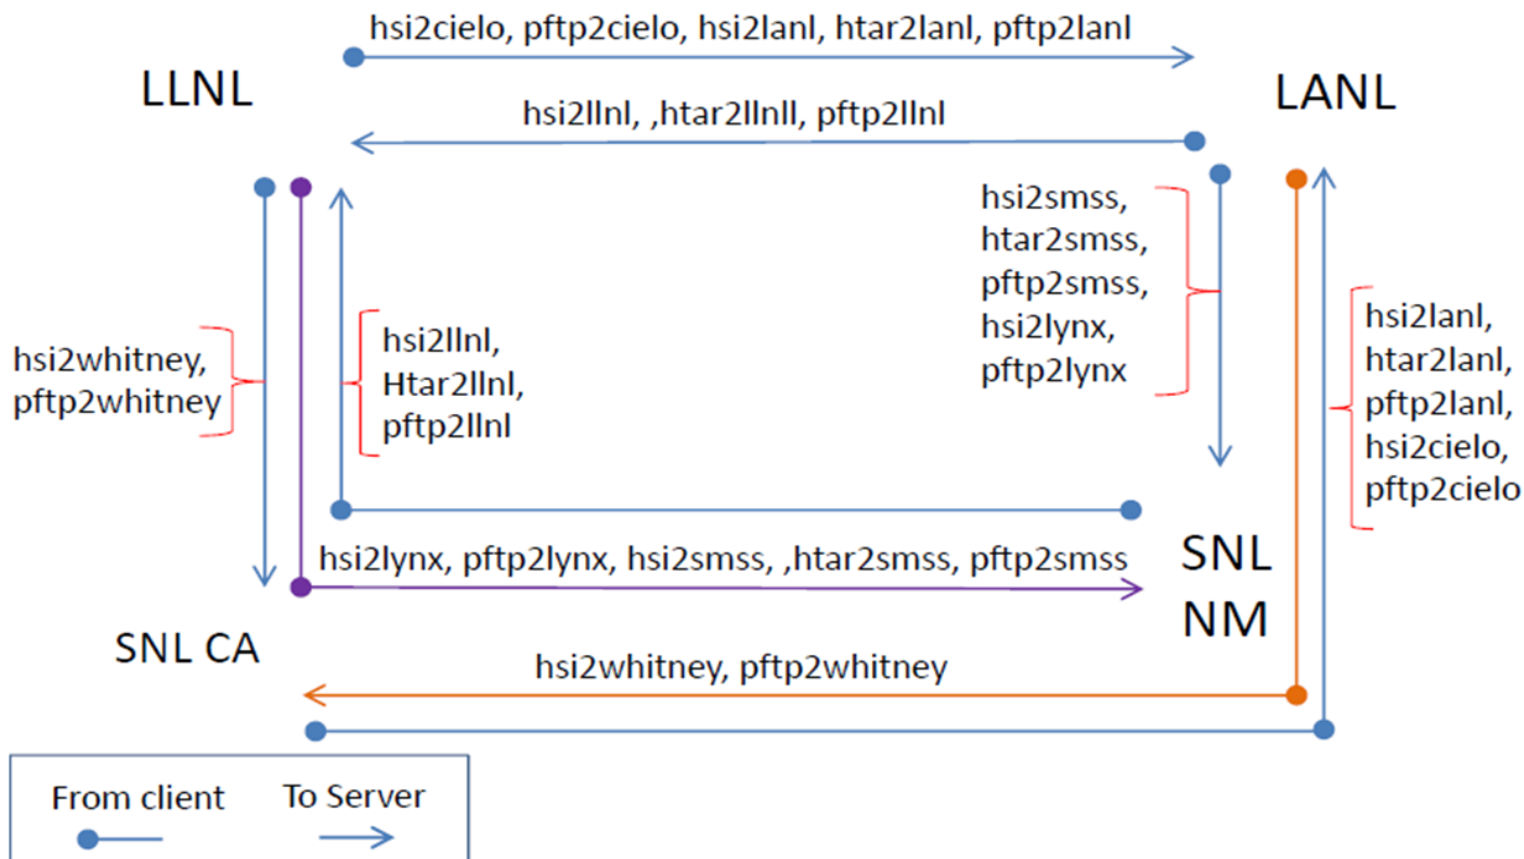
- Primary importance
 - Easy of use – first blush
 - Utility – second...
 - Performance
- Not just a list of features
 - Before, after...

Previous Paths & Destinations



Current Paths & Scripts

Tri-lab Transfer



Questions?

